

SF-sketch: A Two-stage Sketch for Data Streams

Tong Yang¹, Lingtong Liu², Yibo Yan¹, Muhammad Shahzad³, Yulong Shen²

Xiaoming Li¹, Bin Cui¹, Gaogang Xie⁴

¹Peking University, China. ²Xidian University, China.

³North Carolina State University, USA. ⁴ICT, CAS, China

Abstract

A *sketch* is a probabilistic data structure used to record frequencies of items in a multi-set. Sketches are widely used in various fields, especially those that involve processing and storing data streams. In streaming applications with high data rates, a sketch “fills up” very quickly. Thus, its contents are periodically transferred to the remote collector, which is responsible for answering queries. In this paper, we propose a new sketch, called Slim-Fat (SF) sketch, which has a significantly higher accuracy compared to prior art, a much smaller memory footprint, and at the same time achieves the same speed as the best prior sketch. The key idea behind our proposed SF-sketch is to maintain two separate sketches: a small sketch called Slim-subsketch and a large sketch called Fat-subsketch. The Slim-subsketch is periodically transferred to the remote collector for answering queries quickly and accurately. The Fat-subsketch, however, is not transferred to the remote collector because it is used only to assist the Slim-subsketch during the insertions and deletions and is not used to answer queries. We implemented and extensively evaluated SF-sketch along with several prior sketches and compared them side by side. Our experimental results show that SF-sketch outperforms the most widely used CM-sketch by up to 33.1 times in terms of accuracy. We have released the source codes of our proposed sketch as well as existing sketches at Github [2]. The short version of this paper will appear in ICDE 2017 [32].

1 Introduction

1.1 Background and Motivation

A *sketch* is a probabilistic data structure that is used to record frequencies of distinct items in a multi-set. Due to their small memory footprints, high accuracy, and fast speeds of queries, insertions, and deletions, sketches are being extensively used in data stream processing [3], [6], [9], [10], [23], [26], [29], [31], [33]. Sketches are also

being applied in many other fields, such as NLP [15, 17, 18], sparse approximation in compressed sensing [14], and more [12, 16, 20].

Data streams are generated and used in many scenarios, such as network traffic, graph streams, and multimedia streams. For example, a typical application is to measure the number of packets for each flow in the traffic of a network. In most of these applications, the rate at which the data arrives in the stream is very high and a sketch being used to record information in a stream “fills up” very quickly. Consequently, each monitoring node, such as a router or switch, that populates the sketch has to periodically transfer the contents of the filled up sketch to some remote collector [19, 21, 34], which stores the contents of these sketches and answers any queries.

As existing sketches need only a few (*e.g.*, 4) memory accesses for each insertion, the insertion speed of sketches is often fast enough to keep up with the fast rates of data streams. However, the bottleneck lies at the bandwidth and high speed memory of remote controller, because it receives filled up sketches from a large number of monitoring nodes. Consequently, it is critical that the size of the filled up sketch that each monitoring node sends to the collector be very small but still contain enough information that allows the remote collector to answer queries accurately. The accuracy of a sketch in answering queries quantifies how close the value of the frequency estimated from the sketch is to the actual value of the frequency. This paper focuses on the design of a new scheme that not only significantly reduces the size of the sketch that it sends to the collector compared with the existing sketches, but is also more accurate while achieving the same query speed as the best prior sketches.

1.2 Limitations of Prior Art

Here we only introduce four typical sketches: C-sketch, CM-sketch, CU-sketch, and CML-sketch. Charikar *et al.* proposed the Count sketch (C-sketch) [5]. C-sketch experiences two types of errors: over-estimation error,

where the result of a query is a value larger than the true value, and under-estimation error, where the result of a query is a value smaller than the true value. Improving on the C-sketch, Cormode and Muthukrishnan proposed the Count-min (CM) sketch [11], which does not suffer from the under-estimation error, but only from the over-estimation error. And the authors claimed that such one-side error has many benefits [11]. In a further enhancement, Cormode *et al.* proposed the conservative update (CU) sketch [13], which improves the accuracy at the cost of not supporting item deletions. CML-sketch [25] further improves the accuracy at the cost of suffering from both over-estimation and under-estimation errors. Because CM-sketch supports deletions and does not have under-estimation errors, it is still the most popular sketch in practice. As mentioned above, given the desired accuracy and transmission period, smaller sketch leads to less requirement for bandwidth and fast expensive fast memory and faster query speed in the collector.

1.3 Proposed Approach

In this paper, we present a novel two-stage sketch, called Slim-Fat (SF) sketch. The key idea behind our proposed SF-sketch is to maintain two separate sketches: a small sketch called Slim-subsketch and a large sketch called Fat-subsketch. Both subsketches have similar accuracy, and thus, we only need to send the slim-subsketch to the remote collector. Compared to the state-of-the-art, the slim-subsketch achieves significantly higher accuracy and significantly smaller memory footprint while supporting deletions and achieving the same query speed.

Before describing our approach, we first briefly describe how the conventional CM-sketch works because several design choices in SF-sketch are built on it. As shown in Figure 1, a CM-sketch consists of d arrays, where we represent the i^{th} array with A_i . Each array consists of w buckets and each bucket contains one counter. We represent the counter in the j^{th} bucket of the i^{th} array with $A_i[j]$. Each array A_i ($1 \leq i \leq d$), is associated with an independent hash function $h_i(\cdot)$, whose output is uniformly distributed in the range $[1, w]$. In CM-sketch, the initialization operation is simply to set all the $d * w$ counters to zero. To insert an item e , i.e., to increment its frequency stored in the sketch by 1, the CM-sketch computes the d hash functions $h_1(e), h_2(e), \dots, h_d(e)$ and increments the d counters $A_1[h_1(e)], A_2[h_2(e)], \dots, A_d[h_d(e)]$ by 1. For convenience, we call them the d hashed counters. To delete an item e , the CM-sketch computes d hash functions and decrements the d hashed counters by 1. When querying the frequency of an item e , the CM-sketch computes hash functions and returns the value of the smallest one among the d hashed counters. Note that the value returned by the CM-sketch in response to the query for the

frequency of an item will never be smaller than the true value of its frequency. Consequently, CM-sketch suffers only from the under-estimation error, but not from the over-estimation error [11].




Figure 1: The Count-min sketch architecture.

In SF-sketch, the Slim-subsketch, as the name suggests, has significantly fewer counters compared to the Fat-subsketch. The motivation behind keeping the small Slim-subsketch is to use smaller memory while keeping high accuracy. The motivation behind keeping the relatively large Fat-subsketch is to assist the Slim-subsketch during updates so as to make the accuracy of the Slim-subsketch as high as possible. Fat-subsketch uses many more counters compared to Slim-subsketch, but only needs tens of Mega bytes, which is very small compared to the large but cheap memory of current computers. In designing our SF-sketch, we start with a bare bones version of the sketch and make improvements step by step to arrive at its final design. In this process, we present five versions of SF-sketch, namely SF_1 -sketch through SF_4 -sketch, and SF_F -sketch, where each version improves upon the previous versions by addressing some of the limitations in those versions. The version SF_F -sketch is the final design of our SF-sketch.

In SF_1 -sketch, the Slim-subsketch consists of $d \times w$ counters and the Fat-subsketch consists of $d \times w'$ counters, where $w' > w$. The Fat-subsketch in SF_1 -sketch is the same as a standard CM-sketch. The *key idea* behind the design of SF_1 -sketch is that, when inserting an item e , if the value of a counter to which the hash function $h_i(e)$ points is already greater than the real frequency of the item e , then incrementing that counter will only degrade the accuracy. Specifically, when inserting an item e , we first insert it into the Fat-subsketch using the insertion operation of the standard CM-sketch and then query its current frequency from this Fat-subsketch using the query operation of the standard CM-sketch. Suppose the Fat-subsketch estimates e 's current frequency to be c . Next, we compute d hash functions corresponding to the d arrays of the Slim-subsketch, and retrieve d hashed counters. If all d counters are less than or equal to the estimate c , we increment only the smallest counter(s) by 1 in the Slim-subsketch; otherwise, we do nothing. The motivation behind incrementing no counter or only the smallest counter(s) in the Slim-subsketch is twofold. First, by reducing the number of counters that we increment, the over-estimation error reduces. Second, as fewer counters are incremented, their sizes can be reduced,

which reduces the footprint of the Slim-subsketch. Unfortunately, the SF₁-sketch does not support deletions. To well support deletion, we propose several optimization techniques, and finally arrive at the final version – SF_F sketch.

1.4 Technical Challenges

The first technical challenge is to achieve a significantly higher accuracy compared to the CM-sketch, which is currently the most widely used sketch. To address this challenge, we leverage our novel insight that if we reduce the number of counters that are incremented for each insertion, the accuracy will be improved because the extent of over-estimation error will decrease. When inserting a new item, our proposed sketch does not always increment d counters in the Slim-subsketch, *rather increments no counter or only the smallest counter(s)* to avoid under-estimation error. Note that the query is only processed based on the information stored in the Slim-subsketch, which is why we focus on minimizing the number of counter increments per insertion only in the Slim-subsketch. To determine exactly which counters to increment in the Slim-subsketch, our SF-sketch makes use of the Fat-subsketch, which enables it to estimate the number of times the item has already been inserted. SF-sketch then either increments only the smallest counters in the Slim-subsketch if the value of the smallest counters is less than this estimated value, or increments no counter at all.

The second technical challenge is to enable Slim-subsketch to support deletions. It is very difficult to achieve accurate deletions from the Slim-subsketch because to support deletions, one needs to keep track of exactly which counters were incremented when each item was inserted. This information is required to identify the appropriate counters to decrement when deleting an item and to identify the influence of those decrements on other items. Tracking such information is very expensive, both in terms of memory overhead and computational cost. To address this challenge, instead of achieving accurate deletions, *i.e.*, decrementing all those counters that were incremented at the time of inserting the given item, we achieve approximate deletions, *i.e.*, decrementing as many counters in the Slim-subsketch as possible without causing any under-estimation errors.

1.5 Key Contributions

1. We propose a new sketch, namely the SF-sketch, which has higher accuracy compared to the prior art while supporting deletions and keeping the query speed unchanged.
2. We implemented C-sketch, CM-sketch, CU-sketch, CML-sketch and SF-sketch on GPU and multi-core

CPU platforms. We carried out extensive experiments on these two platforms to evaluate and compare the performance of all these sketches. Experimental results show that SF-sketch outperforms CM-sketch by up to 33.1 times in terms of average relative error.

2 Related Work

The structure of the Count sketch (C-sketch) [5] proposed by Charikar *et al.* is exactly the same as the CM-sketch [11] described earlier except that each array A_i is associated with two hash functions $h_i(\cdot)$ and $\delta_i(\cdot)$. Each hash function $\delta_i(\cdot)$ evaluates to -1 or +1 with equal probability. To insert an item e , for all values of $i \in [1, w]$, C-sketch calculates hash functions $h_i(e)$ and $\delta_i(e)$ and adds $\delta_i(e)$ to the counters $A_i[h_i(e)]$. When querying the frequency of item e , C-sketch reports the median of $\{A_1[h_1(e)] \times \delta_1(e), A_2[h_2(e)] \times \delta_2(e) \dots A_d[h_d(e)] \times \delta_d(e)\}$.

Unfortunately, C-sketch suffers from both over-estimation and under-estimation errors. Therefore, several improvements, which do not suffer from the under-estimation errors, have been proposed such as the CM-sketch [11], CU-sketch [13]¹, and Count-Min-Log (CML) sketch [25]. These three sketches all have d arrays of w counters each. To insert an item e , CU-sketch [13] increments only the smallest counter(s) among the d hashed counters. Although CU-sketch improves the query accuracy significantly, its fundamental limitation is that it does not support deletions, and consequently it has not received as wide acceptance in practice as the CM-sketch. CML-sketch is another variant of the CM-sketch that uses logarithm-based approximate counters instead of linear counters [25]. Instead of incrementing one counter per array per insertion, it decides whether or not to increment the counters each time with logarithmic probabilities. This helps in reducing the number of bits for each counter, which in turn allows the sketch to have more counters in the same amount of memory and thus achieve better accuracy. Unfortunately, CML-sketch suffers from both over-estimation and under-estimation errors, and its final version does NOT support deletions. Thorough statistical analysis of various sketches is provided in [27, 28].

A recent work presented Augment sketch (A-sketch), which is a universal framework that can be applied to many existing sketches, especially to those with low accuracy [26]. A-sketch uses a filter to catch heavy hitters (high-frequency items) earlier, and uses classical sketches (such as CM-sketch and C-sketch) to store and

¹Estan and Varghese proposed the CU-sketch which can be combined with other sketches. For convenience, CU-sketch means CM-CU-sketch when it is combined with CM sketch.

query the rest items. In this way, the accuracy could be improved. However, always keeping the most frequent items in the first filter without incurring additional errors is a challenging issue. Complex design and frequent communications between the two filters are unavoidable, making the implementation complicated. Indeed, A-sketch can be applied to our SF-sketch as well. However, according to our test results, as our SF-sketch is already very accurate, combining A-sketch with SF-sketch brings little increase in accuracy but does bring more complexity.

Another class of data structures that can be used to store frequencies of items are the enhanced Bloom filters, such as Spectral Bloom Filters (SBF) [7], Dynamic Count Filters (DCF) [4], and more [31], which indeed can estimate frequencies of items. SBF replaces each bit in the conventional Bloom filter with a counter [7]. To insert an item, the basic version of SBF simply increments all the counters that the item maps to. When querying the frequency of an item, SBF returns the value of the smallest counter(s) among all the counters to which the hash functions map the item to as the estimate of the frequency of that item in the multiset. DCF extends the concept of SBF while improving the memory efficiency of SBF by using two separate filters [4]. The first filter is comprised of fixed size counters while the size of counters in the second filter is dynamically adjusted. The use of two filters, unfortunately, increases the complexity of DCF, which degrades its query and update speeds.

3 The Slim-Fat Sketch

In this section, we present the details of our SF-sketch. To better explain the intuition at work behind the SF-sketch and to justify the design choices we made in developing the SF-sketch, we will start with a basic version and improve it incrementally to arrive at the final design. For each intermediate version of the SF-sketch that we develop while working our way towards the final design, we will first describe its insertion, query, and deletion operations. After that we will discuss its limitations, which will guide us in making our design choices for the next version. In this process, we will present five different versions of SF-sketch, which we name SF₁-sketch through SF₄-sketch, and finally SF_F-sketch, which is our final design. Each version is developed by studying the limitations of its predecessor version and addressing them.

Rationale: In our slim-fat architecture (shown in Figure 2), there is a set of arrays with fewer counters per array called a Slim-subsketch, and a set of arrays with comparatively more counters per array called a Fat-subsketch. When inserting or deleting an item, we first update the Fat-subsketch, and then update the Slim-subsketch based on the observations we make from the Fat-subsketch.




Figure 2: The Slim-Fat sketch architecture.

The key insight at work behind our proposed scheme is that, *when inserting an item, if the smallest one of the d hashed counters is already bigger than its current true frequency, then incrementing any counter only degrades the accuracy.* As the true accuracy is not easy to obtain using small memory, we use Fat-subsketch to give a relatively accurate estimate of the current true frequency. Next, we start with the first version of our slim-fat sketch, i.e., the SF₁-sketch, and discuss its operations and limitations, which will pave the way towards the design of SF₂-sketch and its subsequent versions. Table 1 summarizes the symbols and abbreviations used in this paper.

Table 1: Symbols & abbreviations used in the paper

Symbol	Description
e	Any item that can be handled by SF _{i} -sketch
d	# of arrays in Slim-subsketch and Fat-subsketch
z	# of counters in each bucket of Fat-subsketch of SF ₄ - and SF _F -sketch
w/w'	# of counters or buckets in each array of Slim- / Fat-subsketch
A_i	the i^{th} array in the Slim-subsketch of SF _{i} -sketch
B_i	the i^{th} array in the Fat-subsketch of SF _{i} -sketch
C_i	the i^{th} array in the Deletion-subsketch used in SF ₂ -sketch
$h_i(e)$	the i^{th} hash function used in Slim- and Deletion-subsketch
$g_i(e)$	the i^{th} hash function used in Fat-subsketch
B_e^{\min}	the minimum value among all counters in $\{B_i 1 \leq i \leq d\}$
$\%$	mod operation

3.1 SF₁: Optimizing Accuracy Using One fat-Subsketch

As shown in Figure 2, SF₁-sketch consists of d arrays in both the Slim-subsketch and the Fat-subsketch. *The Fat-subsketch is exactly a standard CM-sketch with many more counters than the Slim-subsketch.* We represent the i^{th} array in the Slim-subsketch with A_i and in the Fat-subsketch with B_i . Each array in the Slim-subsketch consists of w buckets while each array in the Fat-subsketch consists of w' buckets, where $w' > w$. Furthermore, each bucket in both Slim and Fat-subsketches contains one counter. We represent the counter in the j^{th} bucket of the

i^{th} array in the Slim-subsketch with $A_i[j]$, where $1 \leq i \leq d$ and $1 \leq j \leq w$. Similarly, we represent the counter in the k^{th} bucket of the i^{th} array in the Fat-subsketch with $B_i[k]$, where $1 \leq i \leq d$ and $1 \leq k \leq w'$. Each array \mathbf{A}_i is associated with a uniformly distributed independent hash function $h_i(\cdot)$, where the output of $h_i(\cdot)$ lies in the range $[1, w]$. Similarly, each array \mathbf{B}_i is associated with a uniformly distributed independent hash function $g_i(\cdot)$, where the output of $g_i(\cdot)$ lies in the range $[1, w']$. The structure of the SF₁-sketch is shown in Figure 2. The **initialization** operation for the SF₁-sketch is simply setting all counters $A_i[j]$ and $B_i[k]$ to zero, where $1 \leq i \leq d$, $1 \leq j \leq w$, and $1 \leq k \leq w'$.

Insertion: When inserting an item, the SF₁-sketch first inserts it into the Fat-subsketch, and based on the observations made from the Fat-subsketch, increments appropriate counters in the Slim-subsketch. *The insertion operation in the Fat-subsketch is exactly the same as the conventional CM-sketch.* To insert an item e into the Fat-subsketch, we first compute the d hash functions $g_1(e), g_2(e), \dots, g_d(e)$ and increment the d hashed counters $B_1[g_1(e)], B_2[g_2(e)], \dots, B_d[g_d(e)]$ by 1. After inserting the item, we estimate its current frequency of e by finding the minimum value among the d hashed counters we just incremented and represent it with B_e^{\min} . To insert the item e into the Slim-subsketch, we compute the d hash functions and identify the smallest counter(s) among the d hashed counters $A_1[h_1(e)], A_2[h_2(e)], \dots, A_d[h_d(e)]$. If the value of the smallest counter(s) are not smaller than B_e^{\min} , insertion operation ends. Otherwise, we increment the smallest counter(s) by 1. Note that CU-sketch always increments the smallest counter(s). Thus SF₁-sketch is much more accurate than CU-sketch. In other words, $\forall l \in [1, d]$, SF₁-sketch increment all counters $A_l[h_l(e)]$ by one that satisfy the following two conditions: $A_l[h_l(e)] = \min_{i=1}^d A_i[h_i(e)]$, and $A_l[h_l(e)] < B_e^{\min}$.

Query: When querying the frequency of item e , the SF₁-sketch computes the d hash functions $h_1(e), h_2(e), \dots, h_d(e)$, and returns the value of the smallest counter among $A_1[h_1(e)], A_2[h_2(e)], \dots, A_d[h_d(e)]$ as the result of the query. Note that the query is answered only from the Slim-subsketch.

Deletion: SF₁-sketch does not support deletions.

Advantages and Limitations: The key advantage of the SF₁-sketch is that to answer a query it does not access the Fat-subsketch, but only accesses the Slim-subsketch, which keeps the query speed of this sketch as fast as the conventional CM-sketch. Furthermore, note that during the insertion operation, we either increment no counters or increment only the smallest counter(s) in the Slim-subsketch. The smallest counter in the Fat-subsketch gives the upper bound on the number of times that a given item has already been inserted. This strategy

reduces the number of increments in the Slim-subsketch, which has two advantages. First, it reduces the memory footprint of the Slim-subsketch on the expensive and limited memory. Second, due to fewer increments, the over-estimation error is reduced. Unfortunately, the biggest limitation of the SF₁-sketch is that it does not support deletions from the Slim-subsketch. While the Fat-subsketch assists the Slim-subsketch during insertion operation, it cannot assist in the deletion operation because the numbers of counters per array in the Fat- and Slim-subsketches are not the same. This inability to support deletions from the Slim-subsketch limits the practical usability of the SF₁-sketch. In the next version of our SF-sketch, *i.e.*, the SF₂-sketch, we address this limitation while keeping the advantages of the SF₁-sketch.

3.2 SF₂: Supporting Deletion Deletion-subsketch




Figure 3: An example of the deletion problem.

Difficulties for deletions: It is challenging to achieve accurate deletions in SF₁-sketch because to delete items from the Slim-subsketch of SF₁-sketch, one has to keep track of exactly which counters were incremented when inserting each item. Such tracking is difficult and requires large memory and processing overhead. We explain this with help of an example. As shown in Figure 3, consider a Slim-subsketch that has two arrays and two counters per array, where all counters are initialized to 0. Let us first insert two items e_1 and e_2 and then delete the item e_1 . Furthermore, let e_1 maps to $A_1[1]$ and $A_2[1]$ and e_2 maps to $A_1[1]$ and $A_2[2]$. In inserting e_1 , we increment $A_1[1]$ and $A_2[1]$ both to 1. After that, in inserting e_2 , as the current value of $A_1[1]$ is 1 and $A_2[2]$ is 0, we only increment the smaller of the two, *i.e.*, $A_2[2]$ to 1. At this point, $A_1[1] = 1$, $A_1[2] = 0$, $A_2[1] = 1$, and $A_2[2] = 1$. In deleting e_1 , as e_1 maps to both $A_1[1]$ and $A_2[1]$ and as both were incremented at the time of inserting e_1 , if we decrement them both, the query result of e_2 will be 0, *i.e.*, an under-estimation error occurs, which we do not want in our SF-sketch.

Deletion-subsketch: To support deletions, in addition to one Slim-subsketch and one Fat-subsketch just like in the SF₁-sketch, the SF₂-sketch maintains another sketch, called the Deletion-subsketch. The Deletion-subsketch is essentially a standard CM-sketch. Unlike Fat-subsketch, all the parameters ($d, w, h_i(\cdot)$) of the Deletion-subsketch and the Slim-subsketch are exactly the same. For the Deletion-subsketch, we represent the counter in the j^{th} bucket of the i^{th} array with $C_i[j]$, where $1 \leq i \leq d$ and

$1 \leq j \leq w$. Note that the Fat-subsketch helps in deciding which counters to increment in the Slim-subsketch while inserting an item, whereas the Deletion-subsketch helps in deciding which counters to decrement in the Slim-subsketch when deleting an item. The **initialization** operation for the SF₂-sketch consists of simply setting all counters $A_i[j]$, $B_i[k]$, and $C_i[j]$, to 0 ($1 \leq i \leq d$, $1 \leq j \leq w$, and $1 \leq k \leq w'$).

Insertion: The insertion operation of the SF₂-sketch for the Slim- and Fat-subsketches is exactly the same as that of the SF₁-sketch, except that for the SF₂-sketch, we also add information about the incoming item to the Deletion-subsketch. Specifically, to insert an item e into the Deletion-subsketch, we compute d hash functions and increment the d hashed counters $C_1[h_1(e)]$, $C_2[h_2(e)]$, ..., $C_d[h_d(e)]$ by 1.

Query: The query operation of the SF₂-sketch is exactly the same as the SF₁-sketch.

Deletion: To delete an item e from the SF₂-sketch, we first delete it from the Fat-subsketch by decrementing the d counters $B_1[g_1(e)]$, $B_2[g_2(e)]$, ..., $B_d[g_d(e)]$ by 1 and then delete it from the Deletion-subsketch by decrementing the d counters $C_1[h_1(e)]$, $C_2[h_2(e)]$, ..., $C_d[h_d(e)]$ by 1. Finally, we delete it from the Slim-subsketch. We leverage the fact that before deleting the item from the Deletion-subsketch, *each counter in the Slim-subsketch is always less than or equal to the corresponding counter in the Deletion-subsketch*, because when inserting an item, even if a counter in the Slim-subsketch to which the incoming item maps to is not incremented, the corresponding counter in the Deletion-subsketch is always incremented. To delete the item e from the Slim-subsketch, for each $i \in [1, d]$, we compare $A_i[h_i(e)]$ with $C_i[h_i(e)]$ and decrement $A_i[h_i(e)]$ by 1 only when $A_i[h_i(e)] > C_i[h_i(e)]$.

Advantages and Limitations: The SF₂-sketch is advantageous over the SF₁-sketch because it supports deletions. However, it is not efficient in terms of memory usage and update speed because it has to maintain an additional sketch, the Deletion-subsketch, to support deletions from the Slim-subsketch. In the next version of the SF-sketch, *i.e.*, the SF₃-sketch, we address this limitation while keeping the advantages of both SF₁- and SF₂-sketches.

3.3 SF₃: Combining Fat-subsketch with Deletion-subsketch

In SF₃-sketch, we get rid of the separate Deletion-subsketch, and modify the Fat-subsketch so that, in addition to insertions, it can assist deletions in the Slim-subsketch. The Fat-subsketch in the SF₃-sketch is similar to the Fat-subsketch in the SF₁- and SF₂-sketches. However, in the Fat-subsketch of SF₃-sketch, the number of buckets in each array is given by $w' = z \times w$, where z is

a positive integer. In other words, the Fat-subsketch consumes z times as much memory as the Slim-subsketch. The structure of the Slim-subsketch in the SF₃-sketch is exactly the same as the Slim-subsketches in the SF₁- and SF₂-sketches. However, the hash functions $h_i(\cdot)$, where $1 \leq i \leq d$, associated with the Slim-subsketch are now derived from the hash functions $g_i(\cdot)$, where the output of $g_i(\cdot)$ lies in the range $[1, z \times w]$. More specifically,

$$h_i(\cdot) = (g_i(\cdot) - 1) \% w + 1 \quad (1)$$

Consequently, the value of the hash function $h_i(\cdot)$ always lies in the range $[1, w]$, where w is the number of buckets per array in the Slim-subsketch. Note also that calculating the hash function $h_i(\cdot)$ from the hash function $g_i(\cdot)$ using the equation above essentially *associates* each counter $A_i[j]$ in the Slim-subsketch with z counters $B_i[j]$, $B_i[j+w]$, $B_i[j+2w]$, ..., $B_i[j+(z-1)w]$ in the Fat-subsketch. Every time a counter in the Slim-subsketch is incremented, it is certain that one of its *associated* z counters in the Fat-subsketch is also incremented. This further means that the value of a counter in the Slim-subsketch will always be less than or equal to the sum of values of all its associated counters in the Fat-subsketch.

Insertion: When inserting an item e , the SF₃-sketch first inserts it into the Fat-subsketch. For this we compute the d hash functions $g_1(e)$, $g_2(e)$, ..., $g_d(e)$ and increment the d counters $B_1[g_1(e)]$, $B_2[g_2(e)]$, ..., $B_d[g_d(e)]$ by 1. Next, we find the minimum value among all these d counters and represent it with B_e^{\min} . To insert the item e into the Slim-subsketch, we first compute the d hash functions $h_1(e)$, $h_2(e)$, ..., $h_d(e)$ using Equation 1 and then increment all counters $A_i[h_i(e)]$ that satisfy the conditions $A_l[h_l(e)] < B_e^{\min}$ and $A_l[h_l(e)] = \min_{i=1}^d A_i[h_i(e)]$, where $l \in [1, d]$. Note that if $\min_{i=1}^d A_i[h_i(e)] \geq B_e^{\min}$, we do nothing.

Query: The query operation of SF₃-sketches is exactly the same as that of SF₁- and SF₂-sketches.

Deletion: To delete an item from the SF₃-sketch, we first delete it from the Fat-subsketch and then from the Slim-subsketch. To delete the item e from the Fat-subsketch, we first calculate the d hash functions $g_1(e)$, $g_2(e)$, ..., $g_d(e)$ and then decrement the d counters $B_1[g_1(e)]$, $B_2[g_2(e)]$, ..., $B_d[g_d(e)]$ by 1. To delete the item e from the Slim-subsketch, we leverage the fact stated earlier that before deleting the item from the Fat-subsketch, *the value of a counter in the Slim-subsketch is always less than or equal to the sum of values of all its associated counters in the Fat-subsketch*, because when inserting an item, even if a counter in the Slim-subsketch is not incremented, one of the associated counters in the Fat-subsketch is always incremented. To delete the item e from the Slim-subsketch, after deleting it from the Fat-subsketch, for each $i \in [1, d]$, we compare $A_i[h_i(e)]$ with $\sum_{m=0}^{z-1} B_i[h_i(e) + (m \times w)]$ and decrement $A_i[h_i(e)]$ by 1

if $A_i[h_i(e)] > \sum_{m=0}^{z-1} B_i[h_i(e)] + (m \times w)$. Note that each value of $h_i(e)$ is calculated using Equation (1).

Advantages and Limitations: The advantage of SF₃-sketch over the SF₂-sketch is that it does not have to maintain a separate Deletion-subsketch. Unfortunately, it is not efficient in terms of deletion speed because to delete an item, it needs $d \times z$ memory accesses to add the counters in each array of the Fat-subsketch. In the next version of our SF-sketch, *i.e.*, the SF₄-sketch, we address this limitation while keeping the advantages of all three previous versions of the SF-sketch.

3.4 SF₄: Improving Deletion Speed




Figure 4: The SF₄- and SF_F-sketch architecture.

In SF₄-sketch, we modify the Fat-subsketch so that instead of each bucket having one counter, each bucket has z counters. As shown in Figure 4, in the Fat-subsketch of the SF₄-sketch, we have d arrays with $w' = w$ buckets each, and each bucket now contains z counters instead of one counter. We represent the k^{th} counter in the j^{th} bucket of the i^{th} array in the Fat-subsketch with $B_i[j][k]$, where $1 \leq i \leq d$, $1 \leq j \leq w$, and $1 \leq k \leq z$. Each array B_i in the Fat-subsketch is associated with two uniformly distributed independent hash functions: $h_i(\cdot)$ with output in the range $[1, w]$, which maps an item to a bucket in the i^{th} array, and $f_i(\cdot)$ with output in the range $[1, z]$, which maps an item to a counter inside the bucket $B_i[h_i(\cdot)]$ of the i^{th} array. The Slim-subsketch uses the same hash functions $h_i(\cdot)$ as the Fat-subsketch to map items to buckets. Every time a counter in the Slim-subsketch is incremented, it is certain that one of the counters among the z counters in the corresponding bucket of the Fat-subsketch is also incremented. This means that the value of a counter in the Slim-subsketch will always be less than or equal to the *sum of the values of all counters in the corresponding bucket in the Fat-subsketch*.

Insertion: When inserting an item, the SF₄-sketch first inserts it into the Fat-subsketch, and based on the observations it makes from the Fat-subsketch, increments appropriate counters in the Slim-subsketch. Specifically, to insert an item e into the Fat-subsketch, we first compute d hash functions $h_1(e), h_2(e), \dots, h_d(e)$ and another d hash functions $f_1(e), f_2(e), \dots, f_d(e)$ and increment the d counters $B_1[h_1(e)][f_1(e)], B_2[h_2(e)][f_2(e)], \dots, B_d[h_d(e)][f_d(e)]$ by 1. Next, we find the minimum value

among all counters we just incremented and represent it with B_e^{\min} . To insert the item e into the Slim-subsketch, we identify the counters with the smallest value among the d counters $A_1[h_1(e)], A_2[h_2(e)], \dots, A_d[h_d(e)]$ and increment them by 1 only if their values are less than B_e^{\min} . In other words, we increment all counters $A_l[h_l(e)]$ by one that satisfy the conditions $A_l[h_l(e)] = \min_{i=1}^d A_i[h_i(e)]$ and $A_l[h_l(e)] < B_e^{\min}$, where $l \in [1, d]$. If $\min_{i=1}^d A_i[h_i(e)] \geq B_e^{\min}$, we do nothing.

Query: The query operation of the SF₄-sketch is exactly the same as the SF₁-, SF₂-, and SF₃-sketches.

Deletion: To delete an item from the SF₄-sketch, we first delete it from the Fat-subsketch and then from the Slim-subsketch. To delete the item e from the Fat-subsketch, we first calculate the d hash functions $h_1(e), h_2(e), \dots, h_d(e)$ and another d hash functions $f_1(e), f_2(e), \dots, f_d(e)$ and decrement the d counters $B_1[h_1(e)][f_1(e)], B_2[h_2(e)][f_2(e)], \dots, B_d[h_d(e)][f_d(e)]$ by 1. To delete the item e from the Slim-subsketch, we leverage the fact stated earlier that before deleting the item from the Fat-subsketch, *the value of a counter in the Slim-subsketch is always less than or equal to the sum of values of all counters in the corresponding bucket in the Fat-subsketch*. To delete the item e from the Slim-subsketch, after deleting it from the Fat-subsketch, for each $i \in [1, d]$, we compare $A_i[h_i(e)]$ with $\sum_{k=1}^z B_i[h_i(e)][k]$ and decrement counter $A_i[h_i(e)]$ by 1 if $A_i[h_i(e)] > \sum_{k=1}^z B_i[h_i(e)][k]$. Therefore, one deletion from the Fat-subsketch only needs $d \times z \times b/W$ memory accesses, where b is the number of bits of each counter, W is the size of the machine word, and $b < W$.

Advantages and Limitations: The principles behind the SF₄-sketch and the SF₃-sketch are essentially the same. The advantage SF₄-sketch has over SF₃-sketch is that all counters in the Fat-subsketch corresponding to a counter in the Slim-subsketch are now located in the same bucket. Thus, adding the values of the z counters usually only takes a single memory access. Based on SF₄-sketch, our final version SF_F-sketch aims to minimize the over-estimation error.

3.5 SF_F: Reducing Over-estimation Error (The Final Version)

The structure of the SF_F-sketch is exactly the same as SF₄-sketch. The key idea behind the SF_F-sketch is that in updating the counters in the Slim-subsketch, we keep the value of each counter in the Slim-subsketch always *less than or equal to* the value of the *largest counter in the corresponding bucket* of the Fat-subsketch during insertion and deletion operations. Next, we describe how insertion, deletion, and query operations work in SF_F-sketch followed by an analysis of its error and accuracy.

Insertion: The insertion operation of the SF_F-sketch is exactly the same as the insertion operation of the SF₄-sketch.

Query: The query operation of the SF_F-sketch is exactly the same as the previous versions of the SF-sketch.

Deletion: To delete an item from the SF_F-sketch, we first delete it from the Fat-subsketch and then from the Slim-subsketch. To delete an item e from the Slim-subsketch, we first check the d buckets $B_1[h_1(e)], B_2[h_2(e)] \dots B_d[h_d(e)]$. For each $i \in [1, d]$, if $\max_{k=1}^z B_i[h_i(e)][k]$ changes when deleting item e from the Fat-subsketch, we set $A_i[h_i(e)] = \max_{k=1}^z B_i[h_i(e)][k]$ if $A_i[h_i(e)] > \max_{k=1}^z B_i[h_i(e)][k]$. Otherwise, we leave the value of $A_i[h_i(e)]$ unchanged. The key difference between the deletion operation of SF_F-sketch and SF₄-sketch is that in SF_F-sketch, we compare the value of $A_i[h_i(e)]$ with $\max_{k=1}^z B_i[h_i(e)][k]$ instead of $\sum_{k=1}^z B_i[h_i(e)][k]$, which results in significantly reducing the values of counters in the Slim-subsketch.

Advantages: The key advantage of SF_F-sketch over SF₄-sketch is that during deletion operation, it significantly reduces the counter values in the Slim-subsketch because in SF_F-sketch, we compare the values of counters in the Slim-subsketch with the values of the largest counters in the corresponding buckets of the Fat-subsketch instead of comparing them with the sum of the values of all counters in the corresponding buckets of the Fat-subsketch. This significantly reduces the over-estimation error of SF_F-sketch. Note that SF_F-sketch does not suffer from under-estimation error.

Bound on Over-estimation Error: As a query is entirely answered from the Slim-subsketch, the over-estimation error of SF_F-sketch is actually the over-estimation error of the Slim-subsketch. Therefore, next, we calculate the over-estimation error of the Slim-subsketch of the SF_F-sketch. Let α represent the average number of counters in any given array of the Slim-subsketch that are incremented per insertion. Note that for the standard CM-sketch, the value of α is equal to 1 because in the standard CM-sketch, exactly one counter is incremented in each array when inserting an item. For the Slim-subsketch in the SF_F-sketch, α is less than or equal to 1 because the Fat-subsketch helps in reducing the number of counters that are incremented in the Slim-subsketch per insertion. For any given item e , let $f_{(e)}$ represent its actual frequency and let $\hat{f}_{(e)}$ represent the estimate of its frequency returned by the Slim-subsketch of the SF_F-sketch. Let N represent the total number of insertions of all items into the SF_F-sketch. Let $h_i(\cdot)$ represent the hash function associated with the i^{th} array of the Slim-subsketch, where $1 \leq i \leq d$. Let $X_{i,(e)}[j]$ be the random variable that represents the difference between the actual frequency $f_{(e)}$ of the item e and the value of the j^{th} counter in the i^{th} array, i.e., $X_{i,(e)}[j] = A_i[j] - f_{(e)}$, where

$j = h_i(e)$. Due to hash collisions, multiple items will be mapped by the hash function $h_i(\cdot)$ to the counter j , which increases the value of $A_i[j]$ beyond f_e and results in over-estimation error. As all hash function have uniformly distributed output, $\Pr[h_i(e_1) = h_i(e_2)] = 1/w$. Therefore, the expected value of any counter $A_i[j]$, where $1 \leq i \leq d$ and $1 \leq j \leq w$, is $\alpha N/w$. Let ε and δ be two numbers that are related to d and w as follows: $d = \lceil \ln(1/\delta) \rceil$ and $w = \lceil \exp/\varepsilon \rceil$. The expected value of $X_{i,(e)}[j]$ is given by the following expression.

$$\begin{aligned} E(X_{i,(e)}[j]) &= E(A_i[j] - f_{(e)}) \\ &\leq E(A_i[j]) \\ &= \frac{\alpha N}{w} \leq \frac{\varepsilon \alpha}{\exp} N \end{aligned} \quad (2)$$

Finally, we derive the probabilistic bound on the over-estimation error of the Slim-subsketch of the SF_F-sketch.

$$\begin{aligned} \Pr[\hat{f}_{(e)} \geq f_{(e)} + \varepsilon \alpha N] &= \Pr[\forall i, A_i[j] \geq f_{(e)} + \varepsilon \alpha N] \\ &= (\Pr[A_i[j] - f_{(e)} \geq \varepsilon \alpha N])^d \\ &= (\Pr[X_{i,(e)}[j] \geq \varepsilon \alpha N])^d \end{aligned}$$

Substituting the value of $\varepsilon \alpha N$ from Equation (2) into the right side of the equation above, we get

$$\begin{aligned} \Pr[\hat{f}_{(e)} \geq f_{(e)} + \varepsilon \alpha N] &\leq (\Pr[X_{i,(e)}[j] \geq \exp E(X_{i,(e)}[j])])^d \\ \text{Applying Markov's Inequality, we get} \\ \Pr[\hat{f}_{(e)} \geq f_{(e)} + \varepsilon \alpha N] &\leq \exp^{-d} \\ &\leq \delta \end{aligned}$$

Derivation of Correct Rate: The *Correct Rate* of a sketch is defined as the expected percentage of items in the given multi-set for which the query response of the sketch contains no error.

In deriving the correct rate of SF_F-sketch, we make two assumptions: 1) all hash functions are independent; 2) the Fat-subsketch is large enough to have negligible error. Before deriving the correct rate, we first prove the following theorem.

Theorem 1. *In the Slim-subsketch, the value of any given counter is equal to the frequency of the most frequent item that maps to it.*

Proof. We prove this theorem using mathematical induction on number of insertions, represented by k .

Base Case, $k = 0$: The theorem clearly holds for the base case because with no insertions, the frequency of the most frequent item is currently 0, which is also the value of all counters.

Induction Hypothesis, $k = n$: Suppose the statement of the theorem holds true after n insertions.

Induction Step, $k = n + 1$: Let $n + 1^{\text{st}}$ insertion be of any item e that has previously been inserted a times. Let $\alpha_i(k)$ represent the values of the counter $A_i[h_i(e)]$ after k insertions, where $0 \leq i \leq d - 1$. There are two cases to consider: 1) e was the most frequent item when $k = n$; 2) e was not the most frequent item when $k = n$.

Case 1: If e was the most frequent item when $k = n$, then according to our induction hypotheses, $\alpha_i(n) = a$. After inserting e , it will still be the most frequent item and its frequency increases to $a + 1$. The counter $A_i[h_i(e)]$ will be incremented once. Consequently, we get $\alpha_i(n + 1) = a + 1$. Thus for this case, the theorem statement holds because the value of the counter $A_i[h_i(e)]$ after insertion is still equal to the frequency of the most frequent item, which is e .

Case 2: If e was not the most frequent item when $k = n$, then according to our induction hypotheses, $\alpha_i(n) > a$. After inserting e , it may or may not become the most frequent item. If it becomes the most frequent item, it means that $\alpha_i(n) = a + 1$ and as our SF_F scheme, the counter $A_i[h_i(e)]$ will stay unchanged. Consequently, we get $\alpha_i(n + 1) = \alpha_i(n) = a + 1$. Thus for this case, the theorem statement again holds because the value of the counter $A_i[h_i(e)]$ after insertion is equal to the frequency of the new most frequent item, which is e .

After inserting e , if it does not become the most frequent item, then it means $\alpha_i(n) > a + 1$ and as our SF_F-sketch scheme, the counter $A_i[h_i(e)]$ will stay unchanged. Consequently, $\alpha_i(n + 1) = \alpha_i(n) > a + 1$. Thus, the theorem again holds because the value of the counter $A_i[h_i(e)]$ after insertion is still equal to the frequency of the item that was the most frequent after n insertions. \square

Next, we derive the correct rate of the SF_F-sketch. Let v be the number of distinct items inserted into the slim-subsketch and are represented by e_1, e_2, \dots, e_v . Without loss of generality, let the item e_{l+1} be more frequent than e_l , where $1 \leq l \leq v - 1$. Let X be the random variable representing the number of items hashing into the counter $A_i[h_i(e_l)]$ given the item e_l , where $0 \leq i \leq d - 1$ and $1 \leq l \leq v$. Clearly, $X \sim \text{Binomial}(v - 1, 1/w)$.

From Theorem 1, we conclude that if e_l has the highest frequency among all items that map to the given counter $A_i[h_i(e_l)]$, then the query result for e_l will contain no error. Let \mathbb{E} be the event that e_l has the maximum frequency among x items that map to $A_i[h_i(e_l)]$. The probability $P\{\mathbb{E}\}$ is given by the following equation:

$$P\{\mathbb{E}\} = \binom{l-1}{x-1} / \binom{v-1}{x-1} \quad (\text{where } x \leq l)$$

Let P' represent the probability that the query result for e_l from any given counter contains no error. It is given by:

$$\begin{aligned} P' &= \sum_{x=1}^l P\{\mathbb{E}\} \times P\{X = x\} \\ &= \sum_{x=1}^l \frac{\binom{l-1}{x-1}}{\binom{v-1}{x-1}} \binom{v-1}{x-1} \left(\frac{1}{w}\right)^{x-1} \left(1 - \frac{1}{w}\right)^{v-x} = \left(1 - \frac{1}{w}\right)^{v-l} \end{aligned}$$

As there are d counters, the overall probability that the query result of e_l is correct is given by the following equation.

$$P_{\text{CR}}\{e_l\} = 1 - \left(1 - \left(1 - \frac{1}{w}\right)^{v-l}\right)^d$$

The equality above holds when all v items have different frequencies. If two or more items have equal frequencies, the correct rate increases slightly. Consequently, the expected correct rate Cr of slim-subsketch is bound by:

$$Cr \geq \frac{\sum_{l=1}^v P_{\text{CR}}\{e_l\}}{v} = \frac{\sum_{l=1}^v \left(1 - \left(1 - \left(1 - \frac{1}{w}\right)^{v-l}\right)^d\right)}{v} \quad (3)$$

4 Implementation

In this section, we describe our implementation of the sketches on two different computing platforms namely CPU and GPU. We extensively tested and evaluated SF-sketch and compared its performance with prior sketches on these two platforms. Next, we first describe our implementation on the CPU platform and then describe our implementation on the GPU platform.

4.1 CPU Implementation

Our CPU platform comprised a machine with dual 6-core CPUs (24 threads, Intel Xeon CPU E5-2620 @2 GHz) and 62 GB total system memory. Each CPU has three levels of cache memory: L1, L2, and L3. L1 cache is comprised of two 32KB caches, where one cache acts as the data cache and the other acts as the instruction cache. L2 cache is a single 256KB cache and L3 cache is a single 15MB cache. To evaluate the schemes in different types of settings, our implementations on the CPU platform include both single-thread implementation as well as multi-thread implementation. We used C++ as the programming language. In single-thread implementation, for each sketch, we implemented the entire insertion, deletion, and query process within a single thread. In multi-thread implementation, we run each query in a dedicated thread and process it completely inside that thread, observing near-linear growth in query speed with the increase in the number of threads. We will present the results on query speed in more detail in Section 5.3.2.

4.2 GPU Implementation

As GPUs have seen wide acceptance for high-speed data processing, we implemented our sketches on GPUs as well. For these implementations, we employ the basic architecture of GAMT [22]. More specifically, we evaluated the sketches on GPU platform using CUDA 5.0 architecture. We performed our experiments on a NVIDIA GPU (Tesla C2075, 1147 MHz, 5376 MB device memory, 448 CUDA cores). We implemented our sketches on GPU using two prevalent techniques: batch processing and multi-stream pipelining. Next, we describe our implementations for these two techniques.

4.2.1 Batch Processing

Our system architecture is based on CUDA [24], the well-known parallel computing platform created by NVIDIA. In our implementation, a typical query cycle proceeds in following three steps: (1) copy the incoming queries from the CPU to the GPU, (2) execute the query kernel, and (3) copy the result from the GPU back to the CPU. A kernel in CUDA is a function that is called on CPU but executed on GPU. A query kernel is configured with a series of thread blocks, where each block is comprised of a group of working threads. As GPU chips have hundreds and even thousands of cores, batch processing is needed to accelerate GPU-based implementations. Each batch is first filled with a group of independent queries, and then transferred to and executed on the GPU, *i.e.*, as soon as a query arrives, it is buffered until there are enough queries to fill the current batch of queries before transferring the batch to GPU for processing by the query kernels. Note that in practice, not all the queries are processed simultaneously, but rather GPU’s scheduler decides when to process which query. As CPUs support less parallelism compared to GPUs and the additional memory accesses to the CPUs may deteriorate the batch processing performance of GPUs, in our implementation, all d arrays are stored on the GPU to ensure that the operations to access the arrays are executed completely within the GPU.

4.2.2 Multi-Stream Pipeline

As discussed earlier, batch processing is required to take maximum advantage of the massive parallelization that GPU enables. However, waiting for enough queries to fill a batch before sending the batch to GPU results in unnecessary delays. Furthermore, while a large batch does boost the throughput of the GPU, it increases the waiting time before a batch fills and is transferred to GPU for processing. This means that the query that arrived at the start of the current batch will experience significant latency before it is processed. To resolve this throughput-latency dilemma, we utilize the multi-stream technique

featured in NVIDIA Fermi GPU architecture [22, 30]. A stream, in this context, is a sequence of operations that must be executed in a certain order.

As per CUDA architecture, data transfers and kernel executions within different streams can be concurrent as long as the device supports concurrent operations and the host memories used to exchange data between the CPU and the GPU are page-locked. In this way, when one stream is copying data between the CPU and GPU, another stream can execute query kernels in parallel. As a result, the streams behave as a multi-stage pipeline and reduce the total processing time.

Furthermore, a large batch can be divided into several smaller ones, reducing the average lookup latency while keeping the throughput high. Given a batch of requests and a sequence of active streams, the task mapping should be performed in as balanced way as possible to efficiently use GPU’s parallelism. Let b and n denote the batch size and the number of active streams, respectively. If b is just a multiple of n , the whole batch can be evenly divided. Otherwise, the first $b \% n$ streams may need to perform an extra operation. In our implementation, after dividing the whole batch into multiple smaller batches of approximately identical sizes, we use their offset and size information for task mapping. Each small batch is processed by the specified stream, and all streams are launched one after the other to work as a multi-stage pipeline.

5 Experimental Results

We conducted extensive experiments to evaluate the performance of our SF_F-sketch in terms of *accuracy* and *speed*. Onwards, we will refer to the SF_F-sketch as simply the SF-sketch. For comparison, we also implemented and evaluated the performance of four well known sketches, namely the Count-sketch (C-sketch) [5], the CM-sketch [11] and the CU-sketch [13] and one recently proposed sketch, namely the CML-sketch [25]. *CML-sketch and CU-sketch do NOT support deletions. CML-sketch and Count-sketch suffer from both over-estimation and under-estimation errors.*

5.1 Experimental Setup

Datasets: We use three types of datasets: real world traffic, uniform dataset, and skewed dataset. The real world network traffic trace is captured by the main gateway of our campus, while the uniform and the skewed datasets are generated by the well known YCSB [8]. We keep the skewness of our skewed dataset equal to the default value for YCSB, which is 0.99. We use Memcached [1] to record the real frequency of each item to establish the ground truth.




Figure 5: CDF of relative error (uniform).




Figure 6: Average relative error vs. number of insertions (uniform).




Figure 7: Average relative error vs. number of deletions (uniform).




Figure 8: Increase in error due to deletions (uniform).




Figure 9: CDF of relative error (skewed).




Figure 10: Average relative error vs. number of insertions (skewed).




Figure 11: Average relative error vs. number of deletions (skewed).




Figure 12: Increase in error due to deletions (skewed).

Experimental Comparison: As the memory of the monitor is cheap and large enough, thus we assign the same size of memory for the Slim-subsketch and the state-of-the-art sketches both of which will be transmitted to the collector. For update experiments, we compare them by varying item frequencies and operation size, *i.e.*, the number of insertion and deletion operations.

5.2 Experiments on Accuracy

We use *relative error (RE)* to quantify the accuracy of sketches. Let f_e represent the actual frequency of an item e and let \hat{f}_e represent the estimate of the frequency returned by the sketch, the relative error is defined as the ratio $|\hat{f}_e - f_e|/f_e$. To evaluate accuracy, we used 100K distinct items and fixed parameter setting ($d = 5$, $w = 40000$, $z = 3$). We calculated relative errors for different sketches in three settings: (1) by incrementally increasing the number of insertion operations; (2) by incrementally increasing the number of deletion operations; and (3) by first increasing the number of insertion operations and then deleting the inserted items one by one in reverse order. We performed experiments in these three settings for both uniform and skewed workloads. We also conducted experiments to quantify the effect of system parameters on the performance of the sketches. In all our experiments on accuracy evaluation, we use 100K ($= 100 \times 10^3$) distinct items in total.

5.2.1 Uniform Workload

Relative Error CDF: Our experimental results show that the percentage of items for which the relative error

of our SF-sketch is less than 1% is 74.51%, which is 18.8, 4.3, 2.1 and 1.9 times higher than the corresponding percentages for CML, C, CM and CU-sketches, respectively. Figure 5 reports the empirical cumulative distribution function (CDF) of relative error for the 100K distinct items after a total of 10M ($= 10 \times 10^6$) insertions. Specifically, we first inserted the 100K distinct items for a total of 10M times such that the probability of occurrence for each item was uniformly distributed, and then calculated the relative errors in the estimates of the frequencies of those 100K distinct items. In this way, we got 100K values of relative error for each of the five sketches (CML, C, CM, CU and SF-sketches). We then plotted a CDF using the 100K relative error values for each sketch. We observe from Figure 5 that the CDF of the SF-sketch is not only higher than that of the other four sketches but also ascends sharply near relative error of 0. This indicates that the relative error in the estimate of the frequencies of most items, calculated from the SF-sketch, is very close to 0.

Relative Error vs. # of Insertions: Our experimental results show that the average relative error of SF-sketch is [0.6 to 6.2], [4.0 to 24.0], [4.4 to 33.1], and [1.8 to 3.8] times smaller than the average relative errors of CML, C, CM, and CU-sketches, respectively. Figure 6 plots the average relative errors in the estimate of the frequencies of the 100K distinct items obtained from the five sketches for different number of insertions. We observe from this figure that the average relative errors of the five sketches converge to different fixed values with increasing number

of insertions with our SF-sketch being the most accurate. The converged average relative error of our SF-sketch is 6.2, 24.0, 33.1 and 3.8 times smaller than the converged average relative errors of CML, C, CM, and CU-sketch, respectively.

Relative Error vs. # of Deletions: *Our experimental results show that the average relative error of SF-sketch is [2.1 to 23.8] and [2.4 to 33.0] times smaller than the average relative errors of C and CM-sketches, respectively.* Figure 7 plots the average relative errors in the estimate of the frequencies of the 100K distinct items obtained from the three sketches for different number of deletions. Before starting the deletions, we inserted the 100K items 10M times for each sketch. Note that Figure 7 does not include results for CML and CU-sketches because they do not support deletions.

Increase in Error due to Deletions: *Our experimental results show that our SF-sketch loses some accuracy due to deletions, while C and CM-sketches do not. Despite that, the average relative error of SF-sketch is still [2.1 to 24.1] times lower compared to the C-sketch and [2.4 to 33.4] times lower compared to the CM-sketch.* In this experiment, we first inserted the 100K items 10M times, and then deleted them in reverse order as the insertions. After every 100K insertions, we calculated the average relative error for the 100K distinct items and plotted them in Figure 8. Similarly, after every 100K deletions, we calculated the average relative error for the 100K distinct items and plotted them in Figure 8. The lines with a hollow square/triangle/circle indicate average relative errors calculated after insertion operations, while the lines with solid square/triangle/circle indicate average relative errors calculated after deletion operations. Again, Figure 8 does not include results for CML and CU-sketches because they do not support deletions. We observe from this figure that the lines with a hollow square/triangle for the C and CM-sketches perfectly track the corresponding lines with solid ones, which means that deletion operations do not reduce accuracy in the C and CM-sketches. However, the line with a hollow circle for the SF-sketch lie below the corresponding line with solid circle showing that the deletion operation deteriorates the accuracy of SF-sketch. The reason is that when deleting an item, SF-sketch cannot always decrement all counters of Slim-subsketch that it incremented when inserting that item because decrementing all those counters can lead to underestimation error. Thus, SF-sketch supports deletions at the cost of slightly reduced accuracy after deletions.

5.2.2 Skewed Workload

For skewed workloads, we performed exactly the same experiments as for the uniform workloads. The results from these experiments are shown in Figures 9, 10, 11,

and 12. The trends in these figures for the skewed distribution are similar to what we observed for the uniform distribution. Therefore, for the sake of brevity, next, we concisely report the results without describing again how the experiments were conducted.

Relative Error CDF: *Our experimental results, reported in Figure 9, show that in case of skewed workload, the percentage of items for which the relative error of our SF-sketch is less than 1% is 74.30%, which is 21.3, 4.6, 2.1 and 1.7 times higher than the corresponding percentages for CML, C, CM and CU sketches, respectively.*

Relative Error vs. # of Insertions: *Our experimental results, reported in Figure 10, show that in case of skewed workload, the average relative error of SF-sketch is [0.3, 2.8], [3.2, 12.7], [3.0, 14.8], and [1.5, 2.7] times smaller than the average relative errors of CML, C, CM, and CU sketches, respectively.* The converged average relative error of our SF-sketch is 2.8, 12.7, 14.8 and 2.7 times smaller than the converged average relative errors of CML, C, CM, and CU-sketch, respectively.

Relative Error vs. # of Deletions: *Our experimental results, reported in Figure 11, show that for skewed workload, the average relative error of SF-sketch is [2.1 to 12.7] and [1.9 to 14.8] times smaller than the average relative errors of C and CM-sketches, respectively.*

Increase in Error due to Deletions: *Our experimental results, reported in Figure 12, show that SF-sketch loses some accuracy due to deletions, while C and CM-sketches do not. Despite that, the average relative error of SF-sketch is still [2.1 to 12.7] times lower compared to the C-sketch and [1.9 to 14.8] times lower compared to the CM-sketch.*

5.2.3 Real Traffic

We also used real traffic to evaluate the accuracy of sketches. As real traffic does not have deletion operations, we only show results for relative error CDF and relative error with respect to insertions. We have 10M real traffic instances and regard the traffic with the same destination IP address to belong to the same flow. Using this definition of a flow, there are about 230K flows in our data set, and the size distribution of flows is biased with an expected value of 8.1 and variance of 1606. We set $d = 5$, $w = 300000$, $z = 20$ in this set of experiments.

Relative Error CDF: *Our experimental results, reported in Figure 13, show that in case of real world traffic, after a total of 10M insertions of the 230K distinct items, 99.81% items with our SF-sketch have no error, while the percentages of items with CML, C, CM and CU sketches are 72.26%, 79.28%, 96.21% and 99.06%, respectively.*

Relative Error vs. # of Insertions: *Our experimental results, reported in Figure 14, show that for our real world*




Figure 13: CDF of relative error (real).




Figure 14: Average relative error vs. number of insertions (real).




Figure 15: Accuracy vs. # of buckets per array w .




Figure 16: Accuracy vs. # of arrays d .




Figure 17: The insertion and deletion speed of CM-sketch and SF-sketch.




Figure 18: Throughput vs. # of threads.




Figure 19: GPU Throughput vs. batch size.




Figure 20: GPU latency vs. batch size.

traffic, the average relative error of SF-sketch is [25.4 to 4246.2], [98.0 to 1341.8], [11.8 to 17.3] and [3.5 to 4.8] times smaller than the average relative errors of CML, C, CM, and CU sketches, respectively.

5.2.4 Sketch Parameters

Next, we evaluate the effect of changing the system parameters d (the number of arrays) and w (the number of buckets per array) on the accuracy of the sketches. In each experiment to evaluate the effect of system parameters, we insert the 100K distinct items 10M times.

Accuracy vs. w : Our experimental results show that the CU-sketch requires 1.5 times more memory compared to the SF-sketch to achieve close to 1% average relative error. Figure 15 plots the average relative error by varying the number of buckets per array with d fixed at 5. We observe from this figure that increasing the number of buckets per array reduces the average relative error for each sketch. However, we observe that at 30K buckets per array, the average relative error of our SF-sketch reduces to a very small value of just 0.047. On the contrary, the CU-sketch requires 50K buckets per array to achieve an average relative error of 0.049, but note that it does not support deletions. The CML, C and CM-sketch did not achieve close to 0 average relative errors in our experiment.

Accuracy vs. d : Our experimental results show that our SF-sketch achieve an average relative error of 5.6%

using only 3 arrays whereas the CU-sketch takes 6 arrays to come close to the error of SF-sketch and achieves the average relative error of 7.1%. This shows that compared to CU-sketch, at this error rate, SF-sketch takes only half as much memory. Figure 16 plots the average relative error by varying d with w fixed at 40K. We observe from this figure that using 6 arrays, SF-sketch achieve an average relative error of 1.9%. We also observe that increasing the number of arrays reduces the average relative error for all the sketches.

5.3 Experiments on Update Speed and Query Speed

Next, we evaluate the update speed and query speed of the sketches. For GPU platform, the query speed is the throughput, and we also need to evaluate the latency, which is measured in microseconds and quantifies the time duration between submitting a query and receiving the response. We observed from our experiments that the query speeds of all the five sketches (*i.e.*, CML, C, CM, CU, SF sketches) have little differences. This observation was expected because the query operations of these sketches are almost the same. For this reason, we only present the experimental results of the query speed of SF-sketch and CM-sketch.

5.3.1 Update Speed

Our experimental results show that the insertion and deletion speeds of SF-sketch are 0.56 times and 0.46

times of those of CM-sketch, respectively. Our experimental results show that the CM-sketch achieves the fastest insertion and deletion speeds among CML, C, CU sketches. Thus, we only compare our SF-sketch with CM-sketch in terms of update speed (measured in Mups, million updates per second). Figure 17 plots the insertion and deletion speeds of both SF-sketch with CM-sketch. We conducted experiments using many datasets and observed similar results. Figure 17 shows plots for 10 randomly chosen datasets. We use the same sketch parameters as used in Figure 11. Our results also show that the insertion and deletion speeds of SF-sketch are nearly half of those of CM-sketch. This is because SF-sketch needs to update two subsketches for each update. In this experiment, we only use one CPU core to perform updates. One can achieve much faster update speed using multiple cores or GPU.

5.3.2 Multi-core CPU Platform

Our experimental results show that the SF-sketch experienced a throughput gain of about 650K queries per second per thread up to 24 threads. Figure 18 plots the throughput vs. the number of threads for the SF-sketch. We observe from this figure that SF-sketch achieves a throughput of about 1.34M queries per second with a single thread. For this experiment, we performed 10M queries. Using 24 threads, it achieves a throughput of about 16.3M queries per second. We further observed that increasing the number of threads beyond 24 hardly brought about any improvement because our CPU has 6×2 cores, which support $6 \times 2 \times 2 = 24$ threads. This suggests that the query speed of our SF-sketch increases linearly with the number of CPU cores.

5.3.3 GPU Platform

Our experimental results for three different data sets show that the query speed in GPU increases with the increase in the batch size. As shown in Figure 19, for the batch size of 20K queries, the query speed is around 50 million queries per second (Mqps). With increase in the batch size, such as 64K queries per batch, SF-sketch reaches a query speed higher than 110 Mqps.

Our experimental results for three different data sets show that for SF-sketch, to reduce latency, the batch size of 28K is the most optimal for our experimental setup. Figure 20 shows that the average query latency of SF-sketch is below $410 \mu s$ for batch sizes $\leq 28K$. For batch sizes $\geq 32k$, the latency increases to $511 \sim 584 \mu s$.

6 Conclusion

In this paper, we proposed a new sketch for data streams, namely the SF-sketch, which achieves up to 33.1 times higher accuracy compared to the CM-sketch while keeping the query speed as fast as the CM-sketch. The key idea behind our proposed SF-sketch is to maintain two separate sketches, one is called Slim-subsketch and the other is called Fat-subsketch. We send only the Slim-subsketch to the remote collector to achieve accurate and fast query, while the Fat-subsketch enables the Slim-subsketch to achieve high query accuracy. To evaluate and compare the performance of our proposed SF-sketch, we conducted extensive experiments on multi-core CPU and GPU platforms. Our experimental results show that our SF-sketch significantly outperforms the-state-of-the-art in terms of accuracy.

References

- [1] Memcached - a distributed memory object caching system. <http://memcached.org>.
- [2] Source code of SF sketches. <https://github.com/paper2017/SF-sketch>.
- [3] C. C. Aggarwal and S. Y. Philip. On classification of high-cardinality data streams. In *SDM*, volume 10, pages 802–813. SIAM, 2010.
- [4] J. Aguilar-Saborit, P. Trancoso, V. Muntés-Mulero, and J.-L. Larriba-Pey. Dynamic count filters. *ACM SIGMOD Record*, pages 26–32, 2006.
- [5] M. Charikar, K. Chen, and M. Farach-Colton. Finding frequent items in data streams. In *Automata, Languages and Programming*. Springer, 2002.
- [6] A. Chen, Y. Jin, J. Cao, and L. E. Li. Tracking long duration flows in network traffic. In *Proc. IEEE INFOCOM*, 2010.
- [7] S. Cohen and Y. Matias. Spectral bloom filters. In *Proc. ACM SIGMOD*, pages 241–252, 2003.
- [8] B. F. Cooper, A. Silberstein, and et al. Benchmarking cloud serving systems with YCSB. In *Proc. ACM SOCC*, 2010.
- [9] G. Cormode and M. Garofalakis. Sketching streams through the net: Distributed approximate query tracking. In *Proceedings of the 31st international conference on Very large data bases*, pages 13–24. VLDB Endowment, 2005.
- [10] G. Cormode and M. Hadjieleftheriou. Finding frequent items in data streams. *Proceedings of the VLDB Endowment*, 1(2):1530–1541, 2008.
- [11] G. Cormode and S. Muthukrishnan. An improved data stream summary: the count-min sketch and its applications. *Journal of Algorithms*, 55(1):58–75, 2005.
- [12] B. V. Durme and A. Lall. Streaming pointwise mutual information. In *Advances in Neural Information Processing Systems*, pages 1892–1900, 2009.
- [13] C. Estan and G. Varghese. New directions in traffic measurement and accounting. *ACM SIGCOMM CCR*, 32(4), 2002.
- [14] A. C. Gilbert, M. J. Strauss, J. A. Tropp, and R. Vershynin. One sketch for all: fast algorithms for compressed sensing. In *Proceedings of ACM symposium on Theory of computing*, 2007.
- [15] A. Goyal and H. Daumé III. Approximate scalable bounded space sketch for large data nlp. In *Proc. EMNLP*, 2011.
- [16] A. Goyal and H. Daumé III. Approximate scalable bounded space sketch for large data nlp. In *Proc. EMNLP*, 2011.
- [17] A. Goyal, H. Daumé III, and G. Cormode. Sketch algorithms for estimating point queries in nlp. In *Proc. EMNLP-CoNLL*, pages 1093–1103. Association for Computational Linguistics, 2012.
- [18] A. Goyal, H. Daumé III, and S. Venkatasubramanian. Streaming for large scale nlp: Language modeling. In *Proc. NAACL*. Association for Computational Linguistics, 2009.
- [19] N. Handigol, B. Heller, and et al. I know what your packet did last hop: Using packet histories to troubleshoot networks. In *NSDI*, volume 14, pages 71–85, 2014.
- [20] P. Li, K. W. Church, and T. J. Hastie. One sketch for all: Theory and application of conditional random sampling. In *Advances in Neural Information Processing Systems*, pages 953–960, 2009.
- [21] Y. Li, R. Miao, C. Kim, and M. Yu. Flowradar: a better netflow for data centers. In *Proc. USENIX NSDI*, pages 311–324, 2016.
- [22] Y. Li, D. Zhang, A. X. Liu, and J. Zheng. GAMT: a fast and scalable ip lookup engine for gpu-based software routers. In *Proc. IEEE/ACM ANCS*, 2013.
- [23] Z. Liu, A. Manousis, and et al. One sketch to rule them all: Rethinking network flow monitoring with univmon. In *Proc. ACM SIGCOMM*, 2016.
- [24] NVIDIA Corporation. *NVIDIA CUDA C Best Practices Guide, Version 5.0*, Oct. 2012.
- [25] G. Pitel and G. Fouquier. Count-min-log sketch: Approximately counting with approximate counters. *arXiv preprint arXiv:1502.04885*, 2015.
- [26] P. Roy, A. Khan, and G. Alonso. Augmented sketch: Faster and more accurate stream processing. In *Proc. ACM SIGMOD*, 2016.
- [27] F. Rusu and A. Dobra. Statistical analysis of sketch estimators. In *Proceedings of the 2007 ACM SIGMOD international conference on Management of data*, pages 187–198. ACM, 2007.
- [28] F. Rusu and A. Dobra. Sketches for size of join estimation. *ACM Transactions on Database Systems (TODS)*, 33(3):15, 2008.

- [29] D. Thomas, R. Bordawekar, and et al. On efficient query processing of stream counts on the cell processor. In *Proc. IEEE ICDE*, 2009.
- [30] Y. Wang, Y. Zu, and et al. Wire speed name lookup: A gpu-based approach. In *Proc. USENIX NSDI*, pages 199–212, 2013.
- [31] T. Yang, A. X. Liu, M. Shahzad, Y. Zhong, Q. Fu, Z. Li, G. Xie, and X. Li. A shifting bloom filter framework for set queries. In *Proc. VLDB*, 2016.
- [32] T. Yang, L. Liu, Y. Yan, M. Shahzad, Y. Shen, X. Li, B. Cui, and G. Xie. Sf-sketch: A fast, accurate, and memory efficient data structure to store frequencies of data items. In *Proceedings of the ICDE*, 2017.
- [33] T. Yang, G. Xie, Y. Li, Q. Fu, A. X. Liu, Q. Li, and L. Mathy. Guarantee ip lookup performance with fib explosion. In *ACM Conference on SIGCOMM*, pages 39–50, 2015.
- [34] Y. Zhu, N. Kang, and et al. Packet-level telemetry in large datacenter networks. In *Proc. ACM SIGCOMM*, 2015.

Appendix I: Proof of $A_i[j] \leq \max_{k=1}^z B_i[j][k]$

Here, we use induction to prove that for SF_F-sketch, the value of each counter in the Slim-subsketch is always less than or equal to the value of the largest counter in the corresponding bucket of the Fat-subsketch. Before delving into the proof, we first define some notations. Let $A_i^n[j]$ represent the value of the counter in the j^{th} bucket of the i^{th} array in the Slim-subsketch after n updates, where $1 \leq i \leq d$, $1 \leq j \leq w$, and $n \in \mathbb{N}^0$. Similarly, let $B_i^n[j][k]$ represent the value of the k^{th} counter in the j^{th} bucket of the i^{th} array in the fat-subsketch after n updates, where $1 \leq i \leq d$, $1 \leq j \leq w$, $1 \leq k \leq z$, and $n \in \mathbb{N}^0$. As querying does not change any counter in either Fat-subsketch or Slim-subsketch, we only consider insertion and deletion operations for this proof.

Theorem 2. *Each counter in the Slim-subsketch of the SF_F-sketch is always less than or equal to the largest counter in the corresponding bucket of the Fat-subsketch regardless of the number of insertion and deletion operations performed. Formally, $\forall i \in [1, d], j \in [1, w]$, and $n \in \mathbb{N}^0$,*

$$A_i^n[j] \leq \max_{k=1}^z B_i^n[j][k] \quad (4)$$

Proof. We prove this theorem by induction on n , i.e., the number of updates (insertions and deletions).

Base case: The base case occurs when $n = 0$, i.e., no update has yet been done and all counters in both subsketches are currently 0. Hence, $\forall i \in [1, d]$ and $j \in [1, w]$,

$A_i^0[j] = \max_{k=1}^z B_i^0[j][k] = 0$. Thus, the base case of $n = 0$ satisfies Equation (4).

Induction hypothesis: Let for $n = x$, Equation (4) holds true, i.e., $\forall i \in [1, d], j \in [1, w]$, and $n \leq x$, $A_i^x[j] \leq \max_{k=1}^z B_i^x[j][k]$. Our inductive hypothesis is that if Equation (4) holds true for $\forall i \in [1, d]$ and $j \in [1, w]$ when $n \leq x$, then it also holds true when $n = x + 1$, i.e., $A_i^{x+1}[j] \leq \max_{k=1}^z B_i^{x+1}[j][k]$.

Induction step: To prove that the induction hypothesis is correct, we pick an arbitrary counter $A_i[j]$ in the Slim-subsketch, whose current value after x updates is $A_i^x[j]$. We show that the induction hypothesis is true for this arbitrary counter. As our choice of the counter is arbitrary, this in turn implies that all corresponding counters of the Slim- and Fat-subsketches satisfy the inductive hypothesis.

Let the item that is inserted/deleted in the $x + 1^{\text{st}}$ update is represented by e . There are following two cases to consider when performing the $x + 1^{\text{st}}$ update: $j \neq h_i(e)$ and $j = h_i(e)$. We consider them one by one.

Case 1: $j \neq h_i(e)$ This case implies that in inserting or deleting e , neither the value of the counter $A_i[j]$ is modified nor the value of any counter in the bucket $B_i[j]$ is modified. Consequently, $A_i^{x+1}[j] = A_i^x[j]$ and $\forall k \in [1, z], B_i^{x+1}[j][k] = B_i^x[j][k]$. Thus, $A_i^{x+1}[j] \leq \max_{k=1}^z B_i^{x+1}[j][k]$. The inductive hypothesis holds true for Case 1.

Case 2: $j = h_i(e)$ This case implies that in inserting or deleting e , the value of the counter $A_i[j]$ and the value of $\max_{k=1}^z B_i[j]$ may be modified. We consider the insertion and deletion operations separately.

Insertion: For insertion, there are two cases that we need to consider: when $A_i^x[j] < \max_{k=1}^z B_i^x[j][k]$ and when $A_i^x[j] = \max_{k=1}^z B_i^x[j][k]$.

For the first case of insertion, when $A_i^x[j] < \max_{k=1}^z B_i^x[j][k]$, after incrementing the value of a counter in the bucket $B_i[j]$ for the $x + 1^{\text{st}}$ update, clearly the following inequality holds: $\max_{k=1}^z B_i^x[j][k] \leq \max_{k=1}^z B_i^{x+1}[j][k]$. As the value of $A_i[j]$ can at most increment by 1 during an insertion, the following inequality also holds $A_i^{x+1}[j] \leq \max_{k=1}^z B_i^x[j][k]$. Combining these two inequalities, we get $A_i^{x+1}[j] \leq \max_{k=1}^z B_i^x[j][k] \leq \max_{k=1}^z B_i^{x+1}[j][k]$. Removing the middle term, we get $A_i^{x+1}[j] \leq \max_{k=1}^z B_i^{x+1}[j][k]$. Thus, the inductive hypothesis holds true for this first case of insertion.

For the second case of insertion, when $A_i^x[j] = \max_{k=1}^z B_i^x[j][k]$, we use proof by contradiction to prove that the induction hypothesis is true. Assume that the induction hypothesis is not true for this second case of insertion, which means that the inequality $A_i^{x+1}[j] > \max_{k=1}^z B_i^{x+1}[j][k]$ is true. For this inequality to be true, according to our insertion algorithm of SF_F-sketch, following two conditions must be met at

the same time: $\max_{k=1}^z B_i^{x+1}[j][k] = \max_{k=1}^z B_i^x[j][k]$ and $A_i^{x+1}[j] = A_i^x[j] + 1$. Suppose the first condition is satisfied, i.e., $\max_{k=1}^z B_i^{x+1}[j][k] = \max_{k=1}^z B_i^x[j][k]$, which implies that $B_i^x[j][f_i(e)] < \max_{k=1}^z B_i^x[j][k]$. Furthermore, $B_i^x[j][f_i(e)] = B_i^{x+1}[j][f_i(e)] - 1$. We get

$$\begin{aligned} B_i^{x+1}[j][f_i(e)] - 1 &< \max_{k=1}^z B_i^x[j][k] \\ \Rightarrow B_i^{x+1}[j][f_i(e)] &< \max_{k=1}^z B_i^x[j][k] \end{aligned} \quad (5)$$

Suppose the second condition is also satisfied, i.e., $A_i^{x+1}[j] = A_i^x[j] + 1$, which implies that the counter value $A_i^x[j]$ must be minimum among all counters in the Slim-subsketch to which the item e maps, i.e., $A_i^x[j] = \min_{i'=1}^d A_{i'}^x[h_{i'}(e)]$, and this counter value $A_i^x[j]$ must also be smaller than the smallest value among all counters in the Fat-subsketch to which the item e maps, i.e., $A_i^x[j] < \min_{i'=1}^d B_{i'}^{x+1}[h_{i'}(e)][f_{i'}(e)]$. As $\min_{i'=1}^d B_{i'}^{x+1}[h_{i'}(e)][f_{i'}(e)] \leq B_i^{x+1}[j][f_i(e)]$, we get $A_i^x[j] < B_i^{x+1}[j][f_i(e)]$. Applying the result from Equation (5), we get $A_i^x[j] < \max_{k=1}^z B_i^x[j][k]$. This contradicts with the condition of this second case of insertion, which states that $A_i^x[j] = \max_{k=1}^z B_i^x[j][k]$. Thus, the assumption that the induction hypothesis is not true for this second case of insertion, i.e., $A_i^{x+1}[j] > \max_{k=1}^z B_i^{x+1}[j][k]$ is not correct. This implies that, $A_i^{x+1}[j] \leq \max_{k=1}^z B_i^{x+1}[j][k]$, and the induction hypothesis holds true for this second case of insertion. \square

Deletion: For deletion, there are again two cases that we need to consider just like for insertion: when $A_i^x[j] < \max_{k=1}^z B_i^x[j][k]$ and when $A_i^x[j] = \max_{k=1}^z B_i^x[j][k]$.

For the first case of deletion, when $A_i^x[j] < \max_{k=1}^z B_i^x[j][k]$, as the value of $\max_{k=1}^z B_i^x[j][k]$ can at most decrement by 1 during a deletion, the following inequality holds: $A_i^x[j] \leq \max_{k=1}^z B_i^{x+1}[j][k]$. Similarly, as the value of $A_i[j]$ can at most decrement by 1 during a deletion, the following inequality also holds $A_i^{x+1}[j] \leq A_i^x[j]$. Combining these two inequalities, we get $A_i^{x+1}[j] \leq A_i^x[j] \leq \max_{k=1}^z B_i^{x+1}[j][k]$. Removing the middle term, we get $A_i^{x+1}[j] \leq \max_{k=1}^z B_i^{x+1}[j][k]$. Thus, the inductive hypothesis holds true for this first case of deletion.

For the second case of deletion, when $A_i^x[j] = \max_{k=1}^z B_i^x[j][k]$, we use proof by contradiction to prove that the induction hypothesis is true. Assume that the induction hypothesis is not true for this second case of insertion, which means that the inequality $A_i^{x+1}[j] > \max_{k=1}^z B_i^{x+1}[j][k]$ is true. For this inequality to be true, according to our deletion algorithm of SF_F-sketch, following two conditions must be met at the same time: $\max_{k=1}^z B_i^{x+1}[j][k] = \max_{k=1}^z B_i^x[j][k] - 1$ and $A_i^{x+1}[j] = A_i^x[j]$. Suppose the first condition is satisfied, i.e., $\max_{k=1}^z B_i^{x+1}[j][k] = \max_{k=1}^z B_i^x[j][k] - 1$, which implies the following.

$$\max_{k=1}^z B_i^{x+1}[j][k] < \max_{k=1}^z B_i^x[j][k] \quad (6)$$

Suppose the second condition is also satisfied, i.e., $A_i^{x+1}[j] = A_i^x[j]$, which implies that $A_i^x[j] \leq \max_{k=1}^z B_i^{x+1}[j][k]$. Substituting the value of $\max_{k=1}^z B_i^{x+1}[j][k]$ from Equation (6), we get $A_i^x[j] < \max_{k=1}^z B_i^x[j][k]$. This contradicts with the condition of this second case of insertion, which states that $A_i^x[j] = \max_{k=1}^z B_i^x[j][k]$. Thus, the assumption that the induction hypothesis is not true for this second case of deletion, i.e., $A_i^{x+1}[j] > \max_{k=1}^z B_i^{x+1}[j][k]$ is not correct. This implies that, $A_i^{x+1}[j] \leq \max_{k=1}^z B_i^{x+1}[j][k]$, and the induction hypothesis holds true for this second case of deletion. \square

Appendix II: Proof of No Under-estimation Error

Here, we use induction to prove that SF_F-sketch does not incur any under-estimation error. Before delving into the proof, we first define some notations. Similar to the notations defined in the Appendix 6, let $A_i^n[j]$ represent the value of the counter in the j^{th} bucket of the i^{th} array in the Slim-subsketch after n updates, where $1 \leq i \leq d$, $1 \leq j \leq w$, and $n \in \mathbb{N}^0$. Similarly, let $B_i^n[j][k]$ represent the value of the k^{th} counter in the j^{th} bucket of the i^{th} array in the fat-subsketch after n updates, where $1 \leq i \leq d$, $1 \leq j \leq w$, $1 \leq k \leq z$, and $n \in \mathbb{N}^0$. Furthermore, let $SQ^n(e)$ represent estimate of the frequency of an arbitrary item e calculated from the Slim-subsketch after n updates. Similarly, let $FQ^n(e)$ represent estimate of the frequency of an arbitrary item e calculated from the Fat-subsketch after n updates. As Fat-subsketch of the SF_F-sketch behaves exactly like a CM-sketch for any given item e , we directly use the conclusion from [11] that CM-sketch does not suffer from underestimation error, i.e., for any arbitrary item e and $\forall n \in \mathbb{N}^0, FQ^n(e) \geq f_{(e)}^n$, where $f_{(e)}^n$ represents the actual frequency of item e after n updates. As querying does not change any counter in either Fat-subsketch or Slim-subsketch, we only consider insertion and deletion operations (collectively called updates) in this proof.

Theorem 3. *For updates consisting of insertions and deletions, the SF_F-sketch does not incur underestimation error. Formally, for any arbitrary item e and $\forall n \in \mathbb{N}^0$*

$$SQ^n(e) \geq f_{(e)}^n \quad (7)$$

Proof. We prove this theorem by induction on n , i.e., the number of updates (insertions and deletions).

Base case: The base case occurs when $n = 0$, *i.e.*, no update has yet been done and all counters in both sub-sketches are currently 0. Thus, for any item, the frequency returned by the Slim-subsketch at this point will be 0, *i.e.* for any arbitrary item e , $SQ^0(e) = 0$. As no items have yet been inserted, therefore, $f_{(e)}^0 = 0$. Therefore, $SQ^0(e) = f_{(e)}^0$. Thus, the base case of $k = 0$ satisfies Equation (7).

Induction hypothesis: Let for $n = x$, Equation (7) holds true, *i.e.*, for any arbitrary item e , $SQ^x(e) \geq f_{(e)}^x$. Our inductive hypothesis is that if Equation (7) holds true for any arbitrary item e when $n \leq x$, then it also holds true for that item when $n = x + 1$, *i.e.*, for item e , $SQ^{x+1}(e) \geq f_{(e)}^{x+1}$.

Induction step: We use $n = x + 1$ in the inductive step. The value of $SQ^x(e)$ is given as below.

$$SQ^x(e) = A_{i'}^x[h_{i'}(e)] \quad (8)$$

where $i' \in [1, d]$ such that $A_{i'}^x[h_{i'}(e)] = \min_{i=1}^d A_i^x[h_i(e)]$. Suppose the item that we are going to insert/delete in the $x + 1$ st update is e' , where e' can also be any arbitrary item. We choose to use e' in the induction step instead of e to make the inductive step generic. Note that the inductive step holds when $e' = e$ as well as when $e' \neq e$. Next, we apply the inductive step separately to the two cases: 1) item e' is inserted; 2) item e' is deleted.

Insertion: When inserting the item e' , there are two cases. The first case occurs when $h_{i'}(e) = h_{i'}(e')$ regardless of whether $e = e'$ or $e \neq e'$. The second case occurs when $h_{i'}(e) \neq h_{i'}(e')$. Note that this second case occurs only when $e \neq e'$.

For the first case, where $h_{i'}(e) = h_{i'}(e')$, the counter $A_{i'}[h_{i'}(e)]$ may be incremented if $A_{i'}^x[h_{i'}(e)] < \min_{i=1}^d B_i^{x+1}[h_i(e')][f_i(e')]$. In this case,

$$\begin{aligned} A_{i'}^{x+1}[h_{i'}(e)] &= A_{i'}^x[h_{i'}(e)] + 1 \\ &= SQ^x(e) + 1 \\ &\geq f_{(e)}^x + 1 \end{aligned}$$

If $e \neq e'$, then $f_{(e)}^x + 1 > f_{(e)}^{x+1}$, otherwise $f_{(e)}^x + 1 = f_{(e)}^{x+1}$. Combining these two, we get $f_{(e)}^x + 1 \geq f_{(e)}^{x+1}$. Substitut-

ing this inequality in the equation above, we get

$$A_{i'}^{x+1}[h_{i'}(e)] \geq f_{(e)}^x + 1 \geq f_{(e)}^{x+1}$$

As $A_{i'}^{x+1}[h_{i'}(e)] = SQ^{x+1}(e)$, we get

$$SQ^{x+1}(e) \geq f_{(e)}^{x+1}$$

Thus, the inductive hypothesis holds true for the first case.

For the second case, where $h_{i'}(e) \neq h_{i'}(e')$, the counter $A_{i'}[h_{i'}(e)]$ stays unchanged, *i.e.*, $A_{i'}^{x+1}[h_{i'}(e)] = A_{i'}^x[h_{i'}(e)]$, which means that $SQ^{x+1}(e) = SQ^x(e)$. This further implies that $SQ^{x+1}(e) \geq f_{(e)}^x$. As for this case, $e \neq e'$, $f_{(e)}^x = f_{(e)}^{x+1}$, we get $SQ^{x+1}(e) \geq f_{(e)}^{x+1}$. Thus, the inductive hypothesis holds true for the second case as well.

Deletion: According to the deletion operation of the SF_F-sketch, a deletion after x updates may cause a counter $A_i[j]$ in Slim-subsketch to either stay unchanged or be set to $\max_{k=1}^z B_i^{x+1}[h_i(e)][k]$. For each $i \in [1, d]$, if $h_i(e) = h_i(e')$ and $A_i^x[h_i(e)] > \max_{k=1}^z B_i^{x+1}[h_i(e)][k]$, we have

$$\begin{aligned} A_i^{x+1}[h_i(e)] &= \max_{k=1}^z B_i^{x+1}[h_i(e)][k] \\ &\geq B_i^{x+1}[h_i(e)][f_i(e)] \\ &\geq \min_{i'=1}^d B_{i'}^{x+1}[h_{i'}(e)][f_{i'}(e)] \\ &= FQ^{x+1}(e) \\ &\geq f_{(e)}^{x+1} \end{aligned} \quad (9)$$

Otherwise, we have:

$$\begin{aligned} A_i^{x+1}[h_i(e)] &= A_i^x[h_i(e)] \geq \min_{i'=1}^d A_{i'}^x[h_{i'}(e)] \\ &= SQ^x(e) \geq f_{(e)}^x \geq f_{(e)}^{x+1} \end{aligned} \quad (10)$$

Thus, each counter to which e maps stays larger than the actual frequency of e on deletion after x updates, which in turn implies that $SQ^{x+1}(e) \geq f_{(e)}^{x+1}$ and the inductive hypothesis holds true. \square